

CLAIMS

- 1 1. A distributed network storage system including
2 a plurality of file storage nodes that are distributed on and communicate peer-to-
3 peer over the network; and
4 a storage coordinator that manages file storage on the plurality of files storage
5 nodes, the storage coordinator designating selected groups of nodes within the plurality
6 of file storage nodes and directing the nodes in a given group to communicate peer-to-
7 peer to replicate associated group files.

- 1 2. The storage system of claim 1 further including
2 an administrative console that administers policies set by a system administrator,
3 the administrative console providing policies to the storage coordinator; and
4 the storage coordinator designating the groups in accordance with the policies
5 provided by the administrative console.

- 1 3. The storage system of claim 2 wherein
2 the administrative console sets policies that determine a predetermined minimum
3 number of storage nodes in the group; and
4 the storage coordinator selects at least the predetermined minimum number of
5 nodes for each of the respective groups.

- 1 4. The storage system of claim 2 wherein
2 the administrative console provides to the storage coordinator policies that
3 determine the locations of at least one of the nodes in a group relative to the other nodes
4 in the group, and
5 the storage coordinator selects the nodes for the respective groups in accordance
6 with the applicable policies concerning location.

1 5. The storage system of claim 2 in which the storage coordinator is a member of a
2 cluster of storage coordinators that communicate peer-to-peer and share the managing of
3 the file storage on the plurality of file storage nodes.

1 6. The storage system of claim 2 further including
2 a plurality of remote file storage nodes that communicate peer-to-peer over the
3 network;

4 a remote storage coordinator that manages file storage on the plurality of remote
5 file storage nodes, the storage coordinator designating selected groups of nodes within
6 the plurality of remote file storage nodes and directing the nodes in a given group to
7 communicate peer-to-peer to replicate associated group files; and
8 the administrative console providing policies to the remote storage coordinator
9 that govern the operations of the remote storage coordinator.

1 7. The system of claim 6 wherein the storage coordinator and the remote storage
2 coordinator respectively designate groups that include both file storage nodes and remote
3 file storage nodes.

1 8. The system of claim 2 wherein the storage coordinator manages distributed searches
2 of files stored in the system by designating a selected node from each group to search
3 through the respective group files, the storage coordinator accumulating the search results
4 from the selected nodes and providing the results to a requesting node.

1 9. The system of claim 8 wherein the storage coordinator further designates selected
2 nodes that are near neighbors of the requesting node to supply to the requesting node one
3 or more files requested from the search results.

1 10. The system of claim 2 wherein
2 the administrative console provides policies for access to various types of files;
3 and

4 the storage coordinator creates different types of storage lockers for files
5 associated with different access policies and assigns all or portions of a given storage
6 locker to one or more of the selected groups of nodes.

1 11. The storage system of claim 10 wherein the administrative console provides
2 policies for the different types of lockers, and the storage coordinator manages the
3 respective locker in accordance with the policies.

1 12. The storage system of claim 10 wherein the storage coordinator maintains file
2 allocation tables for the respective lockers.

1 13. The storage system of claim 12 wherein the nodes provide entries for the
2 respective file allocation tables, the entries including links to the physical location of the
3 respective files.

1 14. The storage system of claim 13 wherein a given table entry includes a file Id that
2 consists of a storage group Id, a node Id and an associated message number.

1 15. The storage system of claim 2 wherein the administrative console provides to the
2 storage coordinator replication operation policies that relate to one or more of time
3 constraints, interval constraints, and bandwidth constraints, and
4 the storage coordinator manages the replication operations in accordance with the
5 provided policies.

1 16. The storage system of claim 15 wherein a given node performs the replication
2 operation by segmenting a given file into segments and sending a series of numbered
3 messages that contain file segments to the other nodes in the associated group.

1 17. The storage system of claim 12 wherein the nodes compress the files before
2 replicating the files.

- 1 18. The storage system of claim 17 wherein the nodes encrypt the files before
2 compressing the files.
- 1 19. The system of claim 1 wherein the nodes designated as members of a group
2 reserve storage capacity to store the associated group files.
- 1 20. A distributed network storage system including
2 a network of devices with dedicated storage media, some or all of the devices
3 including unused storage capacity; and
4 a storage coordinator that manages storage of network files on the unused storage
5 capacity of the dedicated storage media, the storage coordinator designating selected
6 groups of nodes within the plurality of nodes and directing the nodes in a given group to
7 communicate peer-to-peer to replicate associated group files and to store the replicated
8 files on the storage media utilizing the otherwise unused storage capacity at the
9 respective nodes.
- 1 21. The storage system of claim 20 further including
2 an administrative console that administers policies set by a system administrator,
3 the administrative console providing policies to the storage coordinator; and
4 the storage coordinator designating the groups in accordance with the policies
5 provided by the administrative console.
- 1 22. The storage system of claim 21 wherein
2 the administrative console sets policies that determine a predetermined minimum
3 number of storage nodes in the group; and
4 the storage coordinator selects at least the predetermined minimum number of
5 nodes for each of the respective groups.
- 1 23. The storage system of claim 21 wherein

2 the administrative console provides to the storage coordinator policies that
3 determine the locations of at least one of the nodes in a group relative to the other nodes
4 in the group, and

5 the storage coordinator selects the nodes for the respective groups in accordance
6 with the applicable policies concerning location.

1 24. The storage system of claim 21 in which the storage coordinator is a member of a
2 cluster of storage coordinators that communicate peer-to-peer and share the managing of
3 the file storage on the plurality of file storage nodes.

1 25. The storage system of claim 21 further including
2 a plurality of remote file storage nodes that communicate peer-to-peer over the
3 network;
4 a remote storage coordinator that manages file storage on the plurality of remote
5 file storage nodes, the storage coordinator designating selected groups of nodes within
6 the plurality of remote file storage nodes and directing the nodes in a given group to
7 communicate peer-to-peer to replicate associated group files; and
8 the administrative console providing policies to the remote storage coordinator
9 that govern the operations of the remote storage coordinator.

1 26. The system of claim 25 wherein the storage coordinator and the remote storage
2 coordinator respectively designate groups that include both file storage nodes and remote
3 file storage nodes.

1 27. The system of claim 21 wherein the storage coordinator manages distributed searches
2 of files stored in the system by designating a selected node from each group to search
3 through the respective group files, the storage coordinator accumulating the search results
4 from the selected nodes and providing the results to a requesting node.

1 28. The system of claim 27 wherein the storage coordinator further designates
2 selected nodes that are near neighbors of the requesting node to supply to the requesting
3 node one or more files requested from the search results.

1 29. The system of claim 21 wherein
2 the administrative console provides policies for access to various types of files;
3 and
4 the storage coordinator creates different types of storage lockers for files
5 associated with different access policies and assigns all or portions of a given storage
6 locker to one or more of the selected groups of nodes.

1 30. The storage system of claim 29 wherein the administrative console provides
2 policies for the different types of lockers, and the storage coordinator manages the
3 respective locker in accordance with the policies.

1 31. The storage system of claim 29 wherein the storage coordinator maintains file
2 allocation tables for the respective lockers.

1 32. The storage system of claim 31 wherein the nodes provide entries for the
2 respective file allocation tables, the entries including links to the physical location of the
3 respective files.

1 33. The storage system of claim 32 wherein a given table entry includes a file Id that
2 consists of a storage group Id, a node Id and an associated message number.

1 34. The storage system of claim 21 wherein the administrative console provides to the
2 storage coordinator replication operation policies that relate to one or more of time
3 constraints, interval constraints, and bandwidth constraints, and
4 the storage coordinator manages the replication operations in accordance with the
5 provided policies.

- 1 35. The storage system of claim 34 wherein a given node performs the replication
- 2 operation by segmenting a given file into segments and sending a series of numbered
- 3 messages that contain file segments to the other nodes in the associated group.

- 1 36. The storage system of claim 31 wherein the nodes compress the files before
- 2 replicating the files.